Communicating the Limitations of AI: The Effect of Message Framing and Ownership on Trust in Artificial Intelligence

Taenyun Kim^a and Hayeon Song^{a*}

^aDepartment of Human-AI Interaction, Sungkyunkwan University, Seoul, Republic of Korea

Abstract

Trust plays an essential role in the interaction between humans and artificial intelligence (AI). To promote trust in AI, information about the AI's performance should be communicated well to the users. Accordingly, this paper investigates how information about AI performance should be presented, focusing on message framing and the ownership of decisions. A 2 (ownership: no ownership vs. ownership) $\times$ 3 (message framing: no information vs. negative information vs. positive information) between-subjects experiment was conducted (N = 120). Participants were asked to choose items to help them survive in the desert, supported by an AI decision. The results showed that participants without decision ownership perceived higher trust than those with decision ownership. Also, trust was perceived to be higher when participants were not given performance information than when they were. The results indicate the importance of carefully communicating with AI. The implications of this study are discussed.

Keywords: AI; trust; artificial intelligence; attribute framing; endowment effect

Introduction

Trust is essential in the interaction between humans and Artificial Intelligence (AI; Hengstler, Enkel, & Duelli, 2016). Many people embrace new technologies primarily because of trust (Gefen, Karahanna, & Straub, 2003; Rheu et al., 2021). In a meta-analysis of research on the technology acceptance model (TAM), trust was shown to lead users to perceive that the technology is useful and easy to use and to encourage a positive attitude toward the item and a willingness to use it (Wu et al., 2011). If users lack trust in the system, they may resist the suggestions of the AI assistant (Longoni, Bonezzi, & Morewedge, 2019; Lui & Lamb, 2018) and avoid using the system even if it is beneficial (Gefen et al., 2003; Lee & See, 2004; Li, Hess, & Valacich, 2008; Siau, Sheng, Nah, & Davis, 2004). Trust is especially important when AI is an information source or decision aid (Fogg & Tseng, 1999).

Trust in AI is becoming more important because a growing number of AI applications provide suggestions to help individuals make decisions ranging from low to high risk. For example, Netflix, an online media service provider, uses a recommendation algorithm to help viewers decide what content to watch based on their preferences (Gomez-Urbe & Hunt, 2016). AI is also used to provide suggestions to help make important decisions, such as in healthcare and finance. AI can now diagnose breast cancer more

effectively than radiologists with fewer mistakes and missed cases (McKinney et al., 2020). It can also distinguish skin cancer with expert-level accuracy using mobile applications (de Carvalho Delgado et al., 2019). In finance, AI decision assistants, so-called robo-advisors, are widely used to provide advice on services such as bank accounts and investments (Lui & Lamb, 2018). However, trust in AI is not earned easily. For example, patients trust the medical decision of an AI healthcare provider less than the decision provided by the human doctor even when the AI has comparable performance to the human (Yokoi et al., 2021).

Thus, it is important to provide performance information in a sophisticated manner to promote trust (Bunt, McGrenere, & Conati, 2007; Pu & Chen, 2006). Trust can be fostered when AI explains and demonstrates information about its operations and limitations (Gunning, 2017). In addition, the recent legislation on the “right to explanation” emphasizes the responsibility to communicate with users about AI limitations and performance (Goodman & Flaxman, 2017). However, providing information on AI performance may not always be helpful because information that reveals the system’s uncertainty or algorithmic limits may harm users’ favorable evaluation of AI (Cai, Jongejan, & Holbrook, 2019; Lim & Dey, 2011; Papenmeier, Englebienne, & Seifert, 2019).

Hence, one important factor to consider when AI provides performance information is the way in which the information is presented. Users’ perceptions of the provided information are susceptible to human cognitive processes and biases (Miller, 2019). Even minor changes in phrasing may have a significant impact on people’s perception. (Plous, 1993).

In this study, we first tested the effect of providing AI’s performance information on building trust and then examined ways in which the information should be presented to enhance trust. Specifically, this study tested the effects of framing and ownership in an AI decision support system by conducting a lab-based experiment. The results indicated that individuals trust the AI more willingly when information about the AI’s performance was not made available than when either a positively or negatively framed statement about the performance was provided. The results also indicated that people without decision ownership trust AI more than those with decision ownership. This study contributes to the literature by explaining the theoretical background behind the findings and presenting guidance on ways to enhance trust in AI with message presentation.

Literature Review

Trust in AI and Providing Performance Information

Trust is one’s readiness to be dependent on the acts of others in the hope that these acts will be beneficial to the trustor (Mayer, Davis, & Schoorman, 1995). Trust in AI can be fostered or damaged depending on what and how information is provided. Some studies suggest that presenting an explanation about the algorithmic model can enhance users’ trust in the system (Bunt et al., 2007; Kizilcec, 2016; Pu & Chen, 2006; Zhang & Curley, 2018). Kizilcec (2016) showed that an informative interface could mitigate the decline in trust when users’ expectations were violated but that too much information negatively influences trust. Other studies suggest that presenting information on AI and its performance is not always helpful (Cai et al., 2019; Kizilcec, 2016; Lim & Dey, 2011; Papenmeier et al., 2019; Stowers et al., 2017). Users sometimes have more trust and insight into AI when how the algorithms work is not explained than when the algorithm is explained (Papenmeier et al., 2019).

In particular, providing information can harm trust when the information reveals uncertainty in the system or its limitations (Cai et al., 2019; Lim & Dey, 2011; Papenmeier et al., 2019). For example, information that indicates that the certainty of the suggestion is low can deter users’ trust (Lim & Dey, 2011). Information that indicates the algorithmic limitation of AI can also cause people to trust the system

less (Cai et al., 2019). It is difficult to explain the underlying algorithm of AI (Ananny & Crawford, 2018; Chen & Haynes, 2016; Mason, 2008) because the processes involved in machine learning or deep learning are often too complicated even for experts to fully understand and explain (James et al., 2013). Even if engineers were to succeed in explaining the operation of the algorithms, laypeople might not understand their complicated explanations (Ananny & Crawford, 2018). For this reason, the information should be provided in a way that can help users understand.

According to Lee and Moray (1992), there are three types of information that affect trust in automation. First, *performance* explains “what the automation does.” Second, *the process* explains “how the automation operates.” Lastly, *the purpose* explains “why the automation was developed.” Performance is a competency, or specialized information about current and previous operations (Lee & See, 2004). It would be the easiest to understand among different types of information because its information is complex than others. Thus, we focused on performance information in this study in an effort to answer the following research question:

RQ1. Does providing information about the performance of AI enhance or harm users’ levels of (a) cognitive and (b) behavioral trust in AI?

How to present the information

There are numerous ways to communicate an AI’s function to its users. Human cognitive processes and biases can influence a user’s perception of the explanation (Miller, 2019). Indeed, even a slight change in wording can affect how people think and behave (Plous, 1993).

Framing Effect

The framing effect is a cognitive bias by which people make choices depending on whether the options are presented positively or negatively (Plous, 1993). Levin and colleagues (1998) distinguished three types of framing: risky choice, attribute, and goal framing. *Risky choice* framing manipulates a set of alternatives with different degrees of risk and the probability of affecting people’s risk preferences. *Attribute framing* employs different expressions of attributes or characteristics of certain events or objects, thus influencing how people appraise an item. Lastly, *goal-framing* sets different wordings for the outcome or potential objective of behavior to affect persuasion. Among these types, attribute framing is the most straightforward, enabling researchers to understand how people perceive the respective groups of framing alternatives (Levin et al., 1998).

In attribute framing, the characteristics of an object or event are manipulated depending on the context (Levin et al., 1998). For instance, the attributes of ground beef can be framed as either positive (e.g., 75% lean) or negative (e.g., 25% fat; Levin & Gaeth, 1988). Although the attributes of these two wordings are essentially the same, the evaluation of favorability with regard to accepting an object or event varies depending on whether the message is positively or negatively framed (Davis & Bobko, 1986; Dunegan, 1993; Levin, Schnittjer, & Thee, 1988). Likewise, AI performance characteristics can be framed positively by saying how accurate it is or negatively by saying how erroneous it is. Suppose that there is an AI that provides eight accurate and two inaccurate suggestions for every ten tries. One can either positively say that the AI’s answers are 80% accurate or negatively that the AI’s answers are 20% erroneous. Notably, the consequences of an event are not important in attribute framing; instead, whether an object is framed positively or negatively is determined by the quality and not the quantity of the attribute. Thus, 20% accurate AI is still a positive framing, even though the overall performance of AI is considered to be flawed, while 20% erroneous AI is still a negative framing, although the overall performance of AI is considered good.

Even experts are not immune to the framing effect. When experts make professional decisions related to their expertise, their behaviors are biased by framing. For instance, medical students were more likely to choose a surgery described with positive framings, such as the survival rate, than negative framings, such as the mortality rate (Marteau, 1989). Likewise, engineers tend to assign more funds to projects whose prior performance was framed positively regarding their success, rather than the equivalent negatively framed percentage of failure (Duchon, Dunegan, & Barton, 1989).

The influence of the framing effect is not only limited to the framed property but also influences other non-framed properties: people lower their evaluation even though these properties are not related to the one being framed. This effect also negatively influences decision making (Wong, 2020). People favor positive framing because the information conveyed via positive labeling highlights favorable associations in memory, whereas negative framing evokes unfavorable associations (Levin & Gaeth, 1988).

We expected people to show a more favorable attitude and behavior toward AI when the performance is presented positively rather than negatively because the effect of the framing extends beyond the framed characteristic. (Levin & Gaeth, 1988). Thus, our first hypothesis is as follows:

H1. Participants will show higher (a) cognitive and (b) behavioral trust when the stated performance is framed positively than negatively.

Ownership of the Message

The ownership of a message can also affect users' compliance and trust. If a suggestion is presented after an individual has already developed the ownership of their own decision, they might be reluctant to change their opinions. The *endowment effect* refers to people's tendency to consider what they already possess to be more valuable than what they do not because they perceive ownership of what they have (Kahneman, Knetsch, & Thaler, 1990). The power of ownership is not limited to tangible objects but can also be expanded to abstract constructs such as choices or decisions (Carmon, Wertenbroch, & Zeelenberg, 2003).

Ownership can be developed when an object is psychologically proximate (Dhar & Wertenbroch, 2000; Hoch & Loewenstein, 1991; Lewin, 1938). For instance, people feel an attachment to their choices as they ruminate on the benefits of each alternative option and develop psychological proximity to them (Dhar & Wertenbroch, 2000). Moreover, previous memories of ownership or the expectation of future utilization can also promote attachment to some options (Strahilevitz & Loewenstein, 1998). In summary, people develop an anticipatory perception of possession of the decision by having ownership during the deliberation as they imagine, simulate, recall, or expect the implications of the choices (Carmon et al., 2003).

There are several ways to operationalize ownership. In some studies, ownership was manipulated by creating an endowment of goods (Dhar & Wertenbroch, 2000; Strahilevitz & Loewenstein, 1998). For example, Strahilevitz and Loewenstein (1998) required some participants to possess a mug to create ownership or allowed them to keep the mug in the first half of the experiment to create ownership history while keeping other participants away from the mug. In other studies, ownership was created by mere changes in wording in the script (Carmon et al., 2003; Kahneman et al., 1990). For example, Kahneman and colleagues (1990) told some participants that they were an "owner" of the token while informing other participants that they were a "buyer" of the token. In this study, ownership was manipulated by changing the wording of instructions.

Individuals would be more hesitant to change their own decisions if the suggestion is regarded as a threat to them. When people perceive that the message threatens their freedom of choice, a motivation to resist the offer can emerge; this phenomenon is called reactance (Brehm & Brehm, 1981). *Reactance* is an unpleasant motivational arousal that occurs when an individual's freedom of choice is threatened or lost

(Brehm, 1966; Brehm & Brehm, 1981). Reactance can be exhibited both behaviorally and psychologically (Brehm & Brehm, 1981). Behaviorally, people show reactance by not choosing a suggestion (Brehm, 1966; Engs & Hanson, 1989). Psychologically, individuals demonstrate reactance by having negative feelings or by denouncing the source of a message (Rains, 2013). Therefore, people would be reluctant to change their own decision to follow an AI's suggestion if it were presented after they had already built ownership of their own choices. We hypothesize the following:

H2. Participants will show a higher level of (a) cognitive and (b) behavioral trust when they do not have ownership of their choices than when they do have ownership of their choices.

Method

Participants

This study recruited 128 participants from a university located in Seoul, South Korea, via the official university website. Most participants were students aged 18 to 51 years ($M = 22.72$, $SD = 4$), and 59% of them ($n = 75$) were women. Eight participants were not included in the analysis because they noticed that the items were predetermined rather than created by AI or failed to follow the experimental instructions. Consequently, 120 participants were included in the data analysis.

Experimental Design

The experiment was a 2 (ownership: no ownership vs. ownership) $\times$ 3 (message framing: no information vs. negatively framed information vs. positively framed information about the AI system) between-subject design. The dependent variables were cognitive and behavioral trust in AI.

Task

The Desert Survival Problem (DSP; Lafferty, Eady, Pond, & Synergistics, 1974) was utilized for the experiment. Participants were told that they were the sole survivor of a deadly plane crash and were stranded in the middle of the desert with nowhere to hide from the sun. They managed to recover 15 items and the AI decision-supporting device before the plane was completely destroyed by fire. They did not know their exact location, and the nearest human habitat could be more than 100 km away. In this imaginary life-at-stake situation, participants with the help of AI had to rank the 15 items depending on the importance of each for surviving in the desert. There were no correct or incorrect answers in this task, but experts' ranking and rationale do exist (Lafferty et al., 1974), and the AI suggestion followed this order. This scenario is widely used in human-computer interaction studies, including the perception of computer personality (Moon & Nass, 1996), human-computer teaming (Nass, Fogg, & Moon, 1996), and the credibility of the computer (Burgoon et al., 2000). The materials were created using Python (Van Rossum & Drake, 1995) and the GUI toolkit library PyQt5 (PyQt, 2012).

Procedure

When the participants arrived in the laboratory, they were seated in front of the computer and informed about the DSP and the experiment. They then commenced the experiment using a computer program. In the task, participants were given two rounds to submit their answers. Ownership was manipulated by providing different instructions for the experiment. After choosing their rankings, the participants in the ownership condition were told to press the "save" button. This was done to create the perception that they did not have decision ownership yet by implying that their decision was tentative and could be altered later.

In contrast, those in the no-ownership condition were instructed to press the “confirm” button. This procedure was designed to allow them to think that they have decision ownership by suggesting that their answer had been established and would not be changed. After the first round, participants were shown one of two messages depending on the ownership condition: “Before you decide the rank of the items, you will now check the suggestion of the AI” (i.e., no ownership condition) or “You have just decided the rank of the items. You will now check the suggestion of the AI” (i.e., ownership condition). See Figure 1.

In the second round, participants were shown their initial decision in the first round, alongside the prearranged suggestion.¹ They were then told that AI created these suggestions. Participants were told that they had to “make decisions based on the suggestion of the AI device” (i.e., no ownership condition) or “make a final decision with the AI” (i.e., ownership condition). Participants were also shown one of three different messages about the accuracy of the AI according to the condition they were in: no information (no message presented), negative framing (“This AI device has a 20% error rate”), and positive framing (“This AI has 80% accuracy”).

After participants finished answering the second round, they were instructed to complete the questionnaire. The participants were compensated with 3000 won (approximately 3 US dollars) and then debriefed after the experiment (Fig. 1). The experiment took approximately 30 minutes to complete (both the task and questionnaire).

Measures

Response time for each round was calibrated. We measured the time from the start of the round until the participants pressed the button to submit their answers. The longer the participants spent in the round, the more cognitive effort they were considered to have put into making decisions (Donders, 1969).

Perceived accuracy of the AI was measured by asking the participants how accurate they considered the AI to be. This measurement was considered necessary because each individual might have different thoughts about the performance of AI despite the performance information presented to them (i.e., 80% accuracy in positive framing, 20% error rate in negative framing, or no message in the no-information condition). In the analysis, perceived accuracy was controlled as a covariate because it varies greatly among people based on their prior familiarity with AI and their level of expertise.

Trust was measured in terms of two aspects: cognitive and behavioral. To measure cognitive trust, we used computer credibility (Fogg & Tseng, 1999).² Computer credibility ($\alpha = 0.94$) is composed of

¹ The 15 prearranged items in order of importance were: cosmetic mirror, plastic raincoat, a quart of water per person, flashlight, pocket knife, red and white parachute, .45 caliber pistol, a pair of sunglasses, compress kit with gauze, compass, sectional air map of the area, a book entitled *Edible Animals of the Desert*, 2 quarts of 180 proof Vodka, 100 salt tablets, and 1 quart of rubbing alcohol.

² In the human-computer interaction domain, *trust* and *credibility* are largely undistinguished from one another and are used interchangeably (Fogg & Tseng, 1999). Fogg and Tseng (1999) pointed out that the notions of *trust* and *credibility* with respect to computers should be separated. In addition, Schroeder, Chiou, and Craig (2021) distinguished *credibility*, which is the perceived quality of someone or something, from *trust*, which has been characterized as an attitude toward someone or something with a belief that an agent will assist an individual in achieving their purpose. Fogg and Tseng (1999) acknowledge that *trust* and *credibility* are interrelated and difficult to distinguish. For instance, the dimensions of trust and credibility appear to have much in common. Early definitions of *credibility* characterize it as a mixture of competence, trustworthiness, and goodwill (McCroskey & Teven, 1999).

perceived trustworthiness and perceived expertise. Perceived trustworthiness ($\alpha = 0.89$) was measured using six indices, such as “Truthful,” “well intended,” and “unbiased.” Perceived expertise ($\alpha = 0.92$) was also measured with six indices, including “knowledgeable,” “competent,” and “powerful.” The answers were ranked on a 7-point Likert scale.

Behavioral trust in AI was measured in terms of compliance. Compliance with the AI suggestion was defined as the degree to which the participants changed their decisions after learning the AI’s suggestion. The root mean squared deviation (RMSD) between the participant’s answer ($\hat{\theta}_i$) and the AI suggestion (θ_i) for the 15 items ($n = 15$) was calculated for each round to calibrate compliance. The exact equation is as follows:

$$RMSD = \sqrt{\frac{\sum_{i=1}^n (\hat{\theta}_i - \theta_i)^2}{n}}$$

A smaller RMSD meant that the participant’s answer was closer to the AI suggestion. Compliance was calculated by subtracting the RMSD of the second round from that of the first round. None of the participants scored an RMSD of 0 in the first round, indicating that the participants’ initial rankings did not fully overlap with those of the AI for any participant. A larger compliance value indicated that the participants aligned their opinions more closely with the AI’s suggestion.

Results

Before testing the research question and hypotheses, a paired *t*-test of RMSD and response time between the first and second rounds was conducted to determine how participants regarded the AI’s suggestion. The results showed that the participants’ RMSDs significantly decreased in the second round ($M = 4.67$, $SD = 1.61$) from the first round ($M = 5.87$, $SD = 0.85$), $t(119) = 10.374$, $p < 0.001$, $d = -0.94$. The results also showed that the response time of the second round ($M = 105.37$, $SD = 59.52$) was significantly less than in the first round ($M = 136.74$, $SD = 58.41$), $t(119) = 5.30$, $p < 0.001$, $d = -0.53$. These results demonstrate that, in general, the participants changed their initial decision, reflecting the AI’s suggestions by using less cognitive effort in the second round.

We also conducted a correlation analysis to clarify the relationship between compliance and the RMSD of the first and second rounds. The results showed that compliance was not correlated with the RMSD of the first round ($r[118] = -0.12$, $p = 0.180$). However, there was a strong negative correlation between behavioral trust and the RMSD of the second round ($r[118] = -0.85$, $p < 0.001$). These results indicate that behavioral trust was primarily affected by the participants’ second ordering rather than their initial decision making.

A multivariate analysis of covariance (MANCOVA) was conducted to investigate the main effect of ownership (no ownership vs. ownership) and message framing (no information vs. negative vs. positive framed information about the AI’s performance) to test the research questions and hypotheses. Multivariate outliers were identified and removed from further analyses when the maximum Mahalanobis distance value was greater than the critical χ^2 value. However, no outliers were identified in the analysis.

In conducting MANCOVA, the perceived accuracy of the AI was included as a covariate. The perceived accuracy of the AI was based on the participants’ subjective perceptions. It was controlled

This view of *credibility* is similar to the factors affecting *trust*, which include ability, integrity, and benevolence (Mayer et al., 1995), or *trust in automation*, which is based on performance, process, and purpose (Lee & See, 2004). Thus, in this study, *credibility* is considered to be a cognitive aspect of *trust*.

because the quality of the model influences the way people perceive AI (Smith-Renner et al., 2020). Because there were no definitive answers in the task, the suggestion of the AI could be considered wrong if the participants strongly believed that their answers were correct. In this case, perceived accuracy would be low.

Indeed, the perceived accuracy of AI differed greatly among individuals, ranging from 5% to 87%. The perceived accuracy of the participants displayed a median of 60% (IQR: 47.5-70). The mean perceived accuracy among the participants was 56.52% ($SD = 22.4$). However, we ran an analysis of variance (ANOVA) to ensure that it actually is an individual factor rather than a factor that differs depending on the conditions. The ANOVA results of the perceived accuracy of AI showed no significant main effect of ownership ($F[1, 114] = 0.00, p = 0.983, \eta^2 = 0.000$), message framing ($F[1, 114] = 0.69, p = 0.503, \eta^2 = 0.012$), or the interaction effect ($F[1, 114] = 0.25, p = 0.777, \eta^2 = 0.004$). Consequently, it was controlled as a covariate.

Effect of the presence of framing on trust: RQ

The RQ asks whether providing information about the performance of AI improves cognitive (RQ1a) and behavioral (RQ1b) trust in AI. To test the RQ, we conducted planned contrast tests to compare the no-information condition to other conditions in which the AI's performance information was provided (both positively and negatively framed information conditions; see Figure 2). In this way, the parts of the RQ could be successfully compared between the no-information condition and the two conditions that provided information.

Individuals reported a significantly higher level of cognitive trust, $F(1, 113) = 9.44, p = 0.003, \eta^2 = 0.034$ when no information about the AI's performance was provided under the no-information condition ($M = 4.08, SD = 0.84$) than in the other conditions ($M = 3.66, SD = 0.84$) where such information was provided. On the other hand, individuals' behavioral trust in the no information ($M = 1.46, SD = 1.40$) and information conditions ($M = 1.06, SD = 1.18$) showed no significant differences, $F(1, 113) = 3.26, p = 0.074, \eta^2 = 0.033$). These results showed that providing no information allowed users to evaluate the AI more favorably in terms of cognitive trust than providing them with negative or positive framing.

Effect of positive and negative framing on trust: H1

H1 hypothesized that positive framing, compared to negative framing, would allow a more favorable perception of AI in terms of cognitive (H1a) and behavioral (H1b) trust. In terms of cognitive trust, the main effect of message framing ($F[2, 113] = 2.52, p = 0.009, \eta^2 = 0.034$) was significant. However, Tukey's post-hoc test showed no significant difference of cognitive trust between the positive- ($M = 3.70, SD = 0.76$) and negative-framing condition ($M = 3.61, SD = 0.84, p = 0.814$). In behavioral trust, there was no significant main effect of message framing ($F[2, 113] = 2.54, p = 0.084, \eta^2 = 0.033$). Overall, positive framing did not make participants trust the AI more than negative framing. Thus, H1 was rejected.

Effect of ownership on trust: H2

H2 hypothesized that users would show a higher level of trust in AI if the suggestion is presented when they do not feel ownership of their choices. MANCOVA was conducted to test the difference between the no-ownership and ownership conditions. The results are shown in Figure 2. The results of the MANCOVA showed that those in the no-ownership condition ($M = 3.94, SD = 0.89$), compared to the ownership condition ($M = 3.65, SD = 0.81$), reported significantly higher cognitive trust ($F[1, 113] = 4.98, p = 0.028, \eta^2 = 0.019$). However, there was no main effect of ownership ($F[2, 113] = 0.90, p = 0.345, \eta^2 = 0.006$) on behavioral trust. Overall, although having no ownership of their selection did not increase their behavioral trust in AI, it led participants to trust the AI more than having ownership of their choice from a cognitive point of view. These findings support H2a but not H2b.

Discussion

The current study investigated the effect of AI performance in terms of general trust in AI and explored ways to present information to improve trust. Trust was measured in two ways. Cognitive trust was measured on a self-reported scale of computer credibility, and behavioral trust was measured by calculating the extent to which participants followed the AI's suggestion. The experimental results show that participants given no information about the AI's performance are more likely to perceive higher levels of trust in AI than those given such information, regardless of how the information was framed. These findings imply that providing performance information about AI should be handled with caution to avoid unfavorable consequences on trust. Regarding the effect of message ownership, the results demonstrated that people are more likely to perceive higher levels of trust in AI if the message is presented when they do not feel ownership of their decision.

Cost of providing information about the performance of AI

Our findings show that people perceived higher levels of trust in the AI's suggestions when they have no information about its performance. This suggests that not communicating its limitations yields a higher level of trust in AI for users. These results support previous findings that providing information can harm user perception (Cai et al., 2019; Lim & Dey, 2011; Papenmeier et al., 2019) but contradict studies with the opposite findings (Bunt et al., 2007; Kizilcec, 2016; Pu & Chen, 2006).

There might be several reasons why the current findings support prior findings that providing information can harm trust. These reasons may also explain why conflicting results have been reported in the literature. First, automation bias may be responsible for these results. Automation bias is people's tendency to ignorantly prefer suggestions from automated decision-making systems (Dzindolet, Peterson, Pomranky, Pierce, & Beck, 2003; Parasuraman & Manzey, 2010). When individuals fall into automation bias, they blindly follow the AI's advice, even though the AI can make a mistake, resulting in a disaster in which human lives and resources are lost (Hertz et al., 2019; Parasuraman & Manzey, 2010; Robinette et al., 2016; Salomons et al., 2018). In our study, participants perceived high levels of trust in the automation even without transparent performance information. This blind trust possibly suggests automation bias.

Second, our experiment provided a simple statement of performance information instead of a wordy and obscure explanation that might be difficult for users to understand. Complicated explanations could harm user satisfaction because they are difficult to understand (Narayanan et al., 2018). We only presented limitation-focused information because we felt that this information was directly related to unfavorable results for which the AI might hold legal responsibility as users' "right to know" largely depends on being informed of the AI's limitations.

However, the 20% algorithmic uncertainty in this study might be deemed a critical limitation, which would have been sufficient to harm trust. Previous studies that focused on system uncertainty or algorithmic limitations found that such information has a detrimental effect on users' trust (Cai et al., 2019; Lim & Dey, 2011; Ha et al., 2020; Papenmeier et al., 2019; Robinette, Howard, & Wagner, 2017). For example, Robinette and colleagues (2017) showed that when people find themselves in a high-risk situation, they lose trust in the computer that reveals its limitations. Furthermore, Ha and colleagues (2020) suggested that people trust AI more when the AI provide no information about itself than when information is provided.

In this study, we focused on the framing effect as a way of providing information. However, the results indicate no significant differences between the two sets of attribute framing. This might be due to the context, in which participants had to imagine that their lives were at stake. Thus, it is possible that attribute framing might influence trust in more casual situations, such as music recommendations.

Future studies are needed to investigate the influence of attribute-framing under other circumstances and compare with our findings. Researchers should also investigate factors other than attribute framing that might result in differences in trust. Apart from attribute framing, several other ways of presenting information on AI might exist. For example, there are other types of framing with different characteristics and mechanisms, specifically risk choice and goal-framing (Levin et al., 1998). Future researchers could investigate whether and how these sets of framing affect trust in both high- and low-risk situations.

Importance of the effect of not having ownership of their choices on cognitive trust in AI

As expected, based on the endowment effect (Kahneman et al., 1990; Kahneman, Knetsch, & Thaler, 1991), we found that people perceived higher levels of trust in AI when they did not feel ownership of their decisions. This indicates that, in order to make users trust in the decision support system, it is crucial to convince them that AI is not trying to threaten their freedom of choice, but rather it is collaborating with the user. If the AI provides feedback or evaluation of a decision one already “owns,” one is likely to perceive a lower level of trust in the AI’s suggestion.

The importance of the perception of collaboration in human-AI teaming is reflected in the current finding that people perceive a lower level of trust in AI if they have already built ownership of their choices when the suggestion is provided. People who perceive AI as forcing them to choose a particular option when the suggestion is provided may show reactance (Brehm, 1966; Brehm & Brehm, 1981). This may result in a reluctance to follow the suggestion (Brehm, 1966; Engs & Hanson, 1989) and give rise to negative feelings about the source of information (Rains, 2013). This might be related to the perception of AI as an authoritative boss who pretends to be superior to humans or teammates who work in a nonhierarchical fashion. Hence, it is vital to help users feel that AI does not infringe on their freedom of choice when assisting with their decisions.

However, the results also imply that the perception of ownership should be carefully manipulated because people need to form a healthy trust relationship with the AI. Both insufficient trust (i.e., undertrust) and excessive trust (i.e., overtrust) can be problematic (de Visser et al., 2020). People tend to undertrust the system when they do not believe in its reliability (Wickens, Hollands, Banbury, & Parasuraman, 2015). If they undertrust the AI, they might not use it or might even dislike it, even though it is advantageous (Gefen et al., 2003; Lee & See, 2004; Li et al., 2008; Siau et al., 2004). In this instance, it may be more appropriate to make users feel that they do not have ownership of their choices when the AI provides suggestions by allowing them to perceive that the system is making the decision jointly with them.

On the other hand, people tend to overtrust AI when the task is too complicated or opaque, such as in the healthcare context (Lyell & Coiera, 2017). Overtrust can often cause severe problems, such as unconditionally following the AI’s suggestion, even though it could be erroneous, which would lead to catastrophe with the possible loss of human life and resources (Hertz et al., 2019; Parasuraman & Manzey, 2010; Robinette et al., 2016; Salomons et al., 2018). For example, a doctor who trusts AI complacently and does not verify the suggested results might make an erroneous diagnosis (Bahner, Hüper, & Manzey, 2008; Parasuraman, Molloy, & Singh, 1993). In such cases, it might be more desirable for users to have ownership of their decisions before introducing AI suggestions. They might then inspect the suggestion more thoroughly because they are less likely to think it is reasonable and trust it less than when they do not have ownership of their thoughts beforehand. Therefore, applying different ownership levels in different contexts and tasks is crucial for enabling users to build a healthier trust relationship with AI.

Importantly, in the current study, the level of ownership and framing were manipulated by varying the wording in the description and instructions of the AI system. The procedure for interacting with the AI did not differ between the conditions. However, these minimal differences in nuance created a vast difference in whether users trusted the AI. These findings again show that even the slightest change in

wording is sufficient to elicit different thoughts and responses (Plous, 1993), so it is essential to understand that different presentations of information can have different consequences when designing an AI system.

Limitations and future studies

This study has several limitations. First, this study did not explicitly measure whether or not participants felt more ownership in the ownership condition and less ownership in the no-ownership condition. Future researchers should examine whether the variations in cognitive trust were indeed caused by the perceived ownership discrepancies across the conditions.

Second, there was only a one-shot interaction between the participants and the AI. Thus, it is possible that although the participants evaluated the AI without more credible information, the message or the interaction was insufficient to ensure compliance. This may explain why there was no significant difference in compliance. According to the MODE model, a change in attitude may lead to behavioral changes when sufficient motivation and opportunities exist (Fazio, 1990). Thus, it would be worthwhile to conduct a longitudinal study to investigate whether what we observed in this study holds as users increasingly interact with AI.

Third, this study only provided performance information about the AI rather than process or purpose information, in Lee and See's (2004) work. However, the importance of explanation is not simply about deciding whether to offer information; it is also about what kind of information to offer (see Chiou & Lee, 2021). For example, users need to understand that AI does not offer random suggestions when an explanation is offered. In addition, AI should explain its behavior and decisions (Gunning, 2017; Siau & Wang, 2018). In this respect, the information provided in our scenario, which merely states that the AI has either 80% accuracy (i.e., positive framing) or a 20% error rate (i.e., negative framing) without a detailed explanation, might have been insufficient to justify its limitations and algorithmic certainty to users.

This insufficient information might also lead to a floor effect of perceived accuracy; more than three-quarters of participants reported that the perceived accuracy of the AI was lower than 80% (or 20% error rate). As the participants were explicitly informed about the 80% accuracy during the study, the lower level of perceived accuracy implies that participants did not feel as much trust in the AI as intended. This low perceived accuracy of the AI might have led to the insignificant results between the two sets of attribute framing. Therefore, future researchers should test more diverse explanations. In this regard, Lee and See's (2004) framework of performance, process, and purpose can help test the content of the information provided about the AI. For example, knowing what outcomes will be demonstrated in certain situations (i.e., performance), how the systems are programmed (i.e., process), and why the AI was developed (i.e., purpose) helps build users' trust.

Lastly, this study has several boundary conditions. Our experimental setting straightforwardly delivered a limitation-focused explanation (in terms of content), which was very easy to understand (in terms of communication style). Therefore, our findings in this study may not apply to other contexts, content, or communication styles in the same way. The inconclusive and insignificant results in previous reports of AI performance may be related to these boundary factors. It is essential to consider these boundary factors when examining how the AI's performance is explained. We hope that other researchers will continue investigating this topic by clearly considering and mentioning these boundary factors and identifying new ones that have not yet been studied.

Conclusion

As the AI-based decision support system grows, it is ever more critical to understand how AI communication may affect users' perceptions and decision-making behavior. Communicating about AI's

limitations is an important issue, as the designers of AI systems are responsible for accurately communicating about the limitations. It is tricky, however, because delivering such information can result in reduced trust. In this study, we first examined whether disclosing AI performance information benefits or harms user trust. We did this to clarify the mixed results in the literature. The findings of this study demonstrate that users perceive lower levels of trust when information about the limitations of AI is presented, even though it is simple and easy to understand. We then looked at how the information should be presented to increase trust. Specifically, we tested the effects of decision ownership and framing and showed that ownership of the decision rather than framing is more influential on the user's perceived trust in the AI. We did this because human cognitive bias and message presentation influence trust. The findings reveal that even a minute change in wording could negatively or positively affect user perception. Thus, it is crucial to understand the importance of human-AI interaction to improve human acceptance of and trust in AI while helping users make the best decision based on the AI's suggestion. The results of this study have theoretical and practical implications that are crucial for achieving this and emphasize the need for future studies; much remains to be investigated regarding the relationship between providing explanations and user trust.

Acknowledgment

This research was supported by the Ministry of Science and ICT (MSIT), Korea, under the ICT Challenge and Advanced Network of HRD (ICAN) program (2020-0-01816), supervised by the Institute of Information & Communications Technology Planning & Evaluation (IITP).

About the Authors

Taenyun Kim is a researcher at the Interaction Science Research Center, Sungkyunkwan University. He received M.S. in interaction science and a double-major of B.A. in Psychology and B.S. in Informatics at Sungkyunkwan University. His research interests include trust in autonomous agents and human cognitive bias in automation systems.

Hayeon Song is a Professor in the Department of Interaction Science at Sungkyunkwan University, South Korea. She received a Ph.D. in communication from University of Southern California. Her primary research interest is to investigate ways to use new media as persuasive and educational vehicle for health promotion.

References

- Ananny, M., & Crawford, K. (2018). Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society*, 20(3), 973–989.
- Bahner, J. E., Hüper, A.-D., & Manzey, D. (2008). Misuse of automated decision aids: Complacency, automation bias and the impact of training experience. *International Journal of Human Computer Studies*, 66(9), 688–699.
- Brehm, J. W. (1966). *A theory of psychological reactance*. Academic Press.
- Brehm, S. S., & Brehm, J. W. (1981). *Psychological reactance: A theory of freedom and control*. Academic Press.
- Bunt, A., McGrenere, J., & Conati, C. (2007, July). Understanding the utility of rationale in a mixed-initiative system for GUI customization. In *International Conference on User Modeling* (pp. 147-156). Springer, Berlin, Heidelberg. Burgoon, J. K., Bonito, J. A., Bengtsson, B., Cederberg, C., Lundeberg,

- M., & Allspach, L. (2000). Interactivity in human-computer interaction: A study of credibility, understanding, and influence. *Computers in Human Behavior*, 16(6), 553–574.
- Cai, C. J., Jongejan, J., & Holbrook, J. (2019). The effects of example-based explanations in a machine learning interface. In *Proceedings of the 24th International Conference on Intelligent User Interfaces* (pp. 258–262).
- Carmon, Z., Wertenbroch, K., & Zeelenberg, M. (2003). Option attachment: When deliberating makes choosing feel like losing. *Journal of Consumer Research*, 30(1), 15–29.
- Chen, P. D., & Haynes, R. M. (2016). Transparency for whom? Impacts of accountability movements for institutional researchers and beyond. *New Directions for Institutional Research*, 2015(166), 11–21.
- Das, T., & Teng, B.-S. (2004). The risk-based view of trust: A conceptual framework. *Journal of Business and Psychology*, 19(1), 85–116.
- Davis, M. A., & Bobko, P. (1986). Contextual effects on escalation processes in public sector decision making. *Organizational Behavior and Human Decision Processes*, 37(1), 121–138.
- de Carvalho Delgado Marques, T., Noels, E., Wakkee, M., Udrea, A., & Nijsten, T. (2019). Development of smartphone apps for skin cancer risk assessment: Progress and promise. *Journal of Medical Internet Research*, 21(7), e13376. 10.2196/13376
- de Visser, E. J., Peeters, M. M., Jung, M. F., Kohn, S., Shaw, T. H., Pak, R., & Neerincx, M. A. (2020). Towards a theory of longitudinal trust calibration in human-robot teams. *International Journal of Social Robotics*, 12(2), 459–478.
- Dhar, R., & Wertenbroch, K. (2000). Consumer choice between hedonic and utilitarian goods. *Journal of Marketing Research*, 37(1), 60–71.
- Donders, F. (1969). On the speed of mental processes. *Acta Psychologica*, 30, 412–431.
- Duchon, D., Dunegan, K. J., & Barton, S. L. (1989). Framing the problem and making decisions: The facts are not enough. *IEEE Transactions on Engineering Management*, 36(1), 25–27.
- Dunegan, K. J. (1993). Framing, cognitive modes, and image theory: Toward an understanding of a glass half full. *Journal of Applied Psychology*, 78(3), 491–503.
- Dzindolet, M. T., Peterson, S. A., Pomranky, R. A., Pierce, L. G., & Beck, H. P. (2003). The role of trust in automation reliance. *International Journal of Human-Computer Studies*, 58(6), 697–718.
- Engs, R., & Hanson, D. J. (1989). Reactance theory: A test with collegiate drinking. *Psychological Reports*, 64(3 suppl), 1083–1086.
- Fazio, R. H. (1990). Multiple processes by which attitudes guide behavior: The mode model as an integrative framework. *Advances in Experimental Social Psychology*, 23, 75–109.
- Fogg, B. J., & Tseng, H. (1999). The elements of computer credibility. In *Proceedings of the Sigchi Conference on Human Factors in Computing Systems* (pp. 80–87).

- Gefen, D., Karahanna, E., & Straub, D. W. (2003). Trust and TAM in online shopping: An integrated model. *MIS Quarterly*, 27(1), 51–90.
- Gomez-Urbe, C. A., & Hunt, N. (2016). The Netflix recommender system: Algorithms, business value, and innovation. *ACM Transactions on Management Information Systems (TMIS)*, 6(4), 13.
- Goodman, B., & Flaxman, S. (2017). European Union regulations on algorithmic decision-making and a “right to explanation”. *AI Magazine*, 38(3), 50–57.
- Gunning, D. (2017). Explainable artificial intelligence (XAI). *Defense Advanced Research Projects Agency (DARPA)*. Retrieved from https://www.darpa.mil/attachments/XAIIndustryDay_Final.pptx
- Ha, T., Kim, S., Seo, D., & Lee, S. (2020). Effects of explanation types and perceived risk on trust in autonomous vehicles. *Transportation Research Part F: Traffic Psychology and Behaviour*, 73, 271–280.
- Hengstler, M., Enkel, E., & Duelli, S. (2016). Applied artificial intelligence and trust—the case of autonomous vehicles and medical assistance devices. *Technological Forecasting and Social Change*, 105, 105–120.
- Hertz, N., Shaw, T., de Visser, E. J., & Wiese, E. (2019). Mixing it up: How mixed groups of humans and machines modulate conformity. *Journal of Cognitive Engineering and Decision Making*, 13(4), 242–257.
- Hoch, S. J., & Loewenstein, G. F. (1991). Time-inconsistent preferences and consumer self-control. *Journal of Consumer Research*, 17(4), 492–507.
- James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). *An introduction to statistical learning* (Vol. 112, p. 18). New York: springer. Kahneman, D., Knetsch, J. L., & Thaler, R. H. (1990). Experimental tests of the endowment effect and the Coase theorem. *Journal of Political Economy*, 98(6), 1325–1348.
- Kahneman, D., Knetsch, J. L., & Thaler, R. H. (1991). Anomalies: The endowment effect, loss aversion, and status quo bias. *Journal of Economic Perspectives*, 5(1), 193–206.
- Kizilcec, R. F. (2016). How much information? Effects of transparency on trust in an algorithmic interface. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (pp. 2390–2395).
- Lafferty, J., Eady, P., Pond, A., & Synergistics, H. (1974). *The desert survival problem: A group decision making experience for examining and increasing individual and team effectiveness: Manual*. Experimental Learning Methods.
- Lahno, B. (2004). Three aspects of interpersonal trust. *Analyse & Kritik*, 26(1), 30–47.
- Lee, J. D., & Moray, N. (1992). Trust, control strategies and allocation of function in human-machine systems. *Ergonomics*, 35(10), 1243–1270.
- Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1), 50–80.

- Levin, I. P., & Gaeth, G. J. (1988). How consumers are affected by the framing of attribute information before and after consuming the product. *Journal of Consumer Research*, 15(3), 374–378.
- Levin, I. P., Schneider, S. L., & Gaeth, G. J. (1998). All frames are not created equal: A typology and critical analysis of framing effects. *Organizational Behavior and Human Decision Processes*, 76(2), 149–188.
- Levin, I. P., Schnittjer, S. K., & Thee, S. L. (1988). Information framing effects in social and personal decisions. *Journal of Experimental Social Psychology*, 24(6), 520–529.
- Lewin, K. (1938). *The conceptual representation and the measurement of psychological forces*. Duke University Press.
- Li, X., Hess, T. J., & Valacich, J. S. (2008). Why do we trust new technology? A study of initial trust formation with organizational information systems. *The Journal of Strategic Information Systems*, 17(1), 39–71.
- Lim, B. Y., & Dey, A. K. (2011). Investigating intelligibility for uncertain context-aware applications. In *Proceedings of the 13th International Conference on Ubiquitous Computing* (pp. 415–424).
- Longoni, C., Bonezzi, A., & Morewedge, C. K. (2019). Resistance to medical artificial intelligence. *Journal of Consumer Research*, 46(4), 629–650.
- Lui, A., & Lamb, G. W. (2018). Artificial intelligence and augmented intelligence collaboration: Regaining trust and confidence in the financial sector. *Information & Communications Technology Law*, 27(3), 267–283.
- Lyell, D., & Coiera, E. (2017). Automation bias and verification complexity: A systematic review. *Journal of the American Medical Informatics Association*, 24(2), 423–431.
- Marteau, T. M. (1989). Framing of information: Its influence upon decisions of doctors and patients. *British Journal of Social Psychology*, 28(1), 89–94.
- Mason, M. (2008). Transparency for whom? Information disclosure and power in global environmental governance. *Global Environmental Politics*, 8(2), 8–13.
- Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust. *Academy of Management Review*, 20(3), 709–734.
- McCroskey, J. C., & Teven, J. J. (1999). Goodwill: A reexamination of the construct and its measurement. *Communications Monographs*, 66(1), 90–103.
- McKinney, S. M., Sieniek, M., Godbole, V., Godwin, J., Antropova, N., Ashrafian, H., ... & Shetty, S. (2020). International evaluation of an AI system for breast cancer screening. *Nature*, 577(7788), 89–94.
- Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*, 267, 1–38.
- Moon, Y., & Nass, C. (1996). How “real” are computer personalities? Psychological responses to personality types in human-computer interaction. *Communication Research*, 23(6), 651–674.

- Narayanan, M., Chen, E., He, J., Kim, B., Gershman, S., & Doshi-Velez, F. (2018). How do humans understand explanations from machine learning systems? An evaluation of the human-interpretability of explanation. *arXiv preprint arXiv:1802.00682*.
- Nass, C., Fogg, B., & Moon, Y. (1996). Can computers be teammates? *International Journal of Human-Computer Studies*, 45(6), 669–678.
- Papenmeier, A., Englebienne, G., & Seifert, C. (2019). How model accuracy and explanation fidelity influence user trust. *arXiv preprint arXiv:1907.12652*.
- Parasuraman, R., & Manzey, D. H. (2010). Complacency and bias in human use of automation: An attentional integration. *Human Factors*, 52(3), 381–410.
- Parasuraman, R., Molloy, R., & Singh, I. L. (1993). Performance consequences of automation-induced ‘complacency’. *International Journal of Aviation Psychology*, 3(1), 1–23.
- Plous, S. (1993). *The psychology of judgment and decision making*. McGraw-Hill.
- Pu, P., & Chen, L. (2006). Trust building with explanation interfaces. In *Proceedings of the 11th International Conference on Intelligent User Interfaces* (pp. 93–100).
- PyQt. (2012). *PyQt reference guide*. <https://doc.bccnsoft.com/docs/PyQt4/>
- Rains, S. A. (2013). The nature of psychological reactance revisited: A meta-analytic review. *Human Communication Research*, 39(1), 47–73.
- Rheu, M., Shin, J. Y., Peng, W., & Huh-Yoo, J. (2021). Systematic review: Trust-building factors and implications for conversational agent design. *International Journal of Human–Computer Interaction*, 37(1), 81–96.
- Robinette, P., Howard, A. M., & Wagner, A. R. (2017). Effect of robot performance on human-robot trust in time-critical situations. *IEEE Transactions on Human-Machine Systems*, 47(4), 425–436.
- Robinette, P., Li, W., Allen, R., Howard, A. M., & Wagner, A. R. (2016). Overtrust of robots in emergency evacuation scenarios. In *2016 ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (pp. 101–108).
- Salomons, N., van der Linden, M., Strohkorb Sebo, S., & Scassellati, B. (2018). Humans conform to robots: Disambiguating trust, truth, and conformity. In *2018 ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (pp. 187–195).
- Siau, K., Sheng, H., Nah, F., & Davis, S. (2004). A qualitative investigation on consumer trust in mobile commerce. *International Journal of Electronic Business*, 2(3), 283–300.
- Siau, K., & Wang, W. (2018). Building trust in artificial intelligence, machine learning, and robotics. *Cutter Business Technology Journal*, 31(2), 47–53.
- Smith-Renner, A., Fan, R., Birchfield, M., Wu, T., Boyd-Graber, J., Weld, D. S., & Findlater, L. (2020, April). No explainability without accountability: An empirical study of explanations and feedback in

- interactive ml. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1-13). Stowers, K., Kasdaglis, N., Rupp, M., Chen, J., Barber, D., & Barnes, M. (2017). Insights into human-agent teaming: Intelligent agent transparency and uncertainty. In J. Chen (Ed.), *Advances in human factors in robots and unmanned systems* (pp. 149–160). Springer.
- Strahilevitz, M. A., & Loewenstein, G. (1998). The effect of ownership history on the valuation of objects. *Journal of Consumer Research*, 25(3), 276–289.
- Van Rossum, G., & Drake Jr, F. L. (1995). *Python tutorial* (Vol. 620). Amsterdam: Centrum voor Wiskunde en Informatica.
- Yokoi, R., Eguchi, Y., Fujita, T., & Nakayachi, K. (2021). Artificial intelligence is trusted less than a doctor in medical treatment decisions: Influence of perceived care and value similarity. *International Journal of Human-Computer Interaction*, 37(10), 981-990.
- Wickens, C. D., Hollands, J. G., Banbury, S., & Parasuraman, R. (2015). *Engineering psychology and human performance*. Psychology Press.
- Wong, R. S. (2020). An alternative explanation for attribute framing and spillover effects in multidimensional supplier evaluation and supplier termination: Focusing on asymmetries in attention. *Decision Sciences*, 52(2). 10.1111/deci.12435
- Wu, K., Zhao, Y., Zhu, Q., Tan, X., & Zheng, H. (2011). A meta-analysis of the impact of trust on technology acceptance model: Investigation of moderating influence of subject and context type. *International Journal of Information Management*, 31(6), 572–581.
- Zhang, J., & Curley, S. P. (2018). Exploring explanation effects on consumers' trust in online recommender agents. *International Journal of Human-Computer Interaction*, 34(5), 421–432.

Figure 1. Overview of the experiment. As participants entered the laboratory, they were introduced to the study. In the first round, they were told to “confirm” or “save” their decisions. They received a priming message before the second round, depending on the ownership condition (no ownership vs. ownership of their own decisions). In the second round, they were shown messages with negative framing (“20% error rate”), positive framing (“80% accuracy”), or no information at all about the performance of the AI depending on the framing condition (negative vs. positive vs. no information) and submitted their answers. After the second round, they completed the post-experiment questionnaire.

Figure 2. Experimental results: (a) Framing and (b) Ownership. Participants perceived higher levels of trust when AI did not provide the performance information than it provided the information regardless of the framing types. Also, those without decision ownership perceived higher levels of trust than those with decision ownership. The error bars indicate standard errors. Significance level: $*p < 0.05$